

MSL Relay Coordination and Tactical Planning in the Era of InSight, MAVEN, and TGO

Rachael Collins, Pegah Pashai, Emma Young, Chris Bennett, Sharon Laubach, Steven Thomas
NASA Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
ipe_comm@jpl.caltech.edu

Abstract—This study examines an approach for optimizing the scheduling of regular relay communications between the Mars Science Laboratory (MSL) rover and the non-sun-synchronous Mars orbiters MAVEN and TGO as well as the impacts of the approaching InSight landing on MSL relay and tactical planning. Rover operations require knowledge of recently executed activities, termed decisional data, in order to inform tactical activity planning. As a result, timely and routine data return is critical for nominal rover operations. Mars orbiters are used as relay assets to achieve such timeliness. They also provide greater overall rover data throughput considering their larger data transfer capacity between Mars and Earth. Relay opportunities and their performance are thus tightly coupled to MSL's operations efficiency and science return.

With InSight landing only 600 kilometers away and at the same longitude as MSL, orbiter view periods will be shared between the missions, resulting in fewer relay opportunities for MSL. The introduction of MAVEN and TGO as relay assets helps to alleviate this, but the orbit geometries of these spacecraft introduce their own challenges. Unlike sun-synchronous orbiters MRO and ODY, the timing of MAVEN and TGO overflights walks sol-to-sol, resulting in seasonal variations that preclude their usability. The overflights may occur too early in the sol to enable science activities or too late in the sol to be decisional for the subsequent planning cycle. Moreover, the highly elliptical orbit of MAVEN results in much longer view periods as well as intervals of lower or higher data volume return.

With the introduction of InSight, MAVEN, and TGO, the MSL mission undertook a design effort in order to define new overflight selection criteria and identify the impact to operational efficiency. Instead of selecting all usable relay opportunities, as was the case with just MRO and ODY, this new paradigm requires deconflicting and down-selecting from available overflights. The overflight selection algorithm presented in this study selects based on key overflight metrics such as timing, the predicted data volume return, and the latency between the relay and data arrival to Earth. The relative priority of each of these metrics are scenario-specific; thus, the algorithm is flexible and configurable for when mission priorities evolve. Additionally, operational constraints and considerations such as human factors are applied. The resulting tactical planning timeline post-InSight landing suggests comparable operational efficiency to the pre-InSight era but yields more variation in the timing of the planning shifts, adding strain on the MSL planning team.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. BACKGROUND	2
3. CHALLENGES WITH THE INTRODUCTION OF INSIGHT, MAVEN, & TGO	3
4. PROCESS AND TOOL IMPLEMENTATION.....	5
5. PLANNING TIMELINE IMPACTS	10
6. CONCLUSIONS.....	11
APPENDIX	11
ACKNOWLEDGMENTS	11
REFERENCES	12

1. INTRODUCTION

Mars Science Laboratory (MSL) landed on Mars in August 2012. Since that time, the Curiosity rover has been exploring Gale Crater in an effort to better understand the history and habitability of Mars. With 11 science instruments, Curiosity generates significant amounts of data during each sol (Martian day) of operations. To relay this data back to Earth, MSL downlinks data to orbiters in the Mars Relay Network via an Electra-Lite UHF radio. These orbiters then use larger telecom systems to downlink rover data back to Earth at faster data rates. MSL data is then analyzed by the engineers and science team and is used to enable planning for the next sol of operations on Mars. The data to enable planning of future sols is referred to as decisional data and is a primary driver for MSL to maintain operations efficiency. Historically, MSL has relied on the sun-synchronous orbiters at Mars, including Mars Odyssey (ODY) and Mars Reconnaissance Orbiter (MRO), to downlink this data. MSL's mission was designed so that the consistent overflight timing of these orbiters in the Martian day enables the return of MSL's decisional data.

InSight arrives at Mars in November 2018 and lands in close proximity to MSL, resulting in shared Mars orbiter view periods between the two landers. The relay support of MRO and ODY is not sufficient to meet both MSL and InSight's objectives. Because of this, MAVEN and the European Space Agency's ExoMars Trace Gas Orbiter (TGO) were added as relay assets to the Mars Relay Network. However, the inclusion of these new spacecraft present operational challenges for MSL due to their orbit geometries and planning cadences. MSL had to work with the four missions providing relay support (MRO, ODY, TGO, MAVEN) as well as the InSight mission to design and develop a solution to ensure that all missions would be able to satisfy their respective mission objectives to explore Mars. The solution presented aims to maximize MSL's operations efficiency, given the challenges and constraints of the new paradigm.

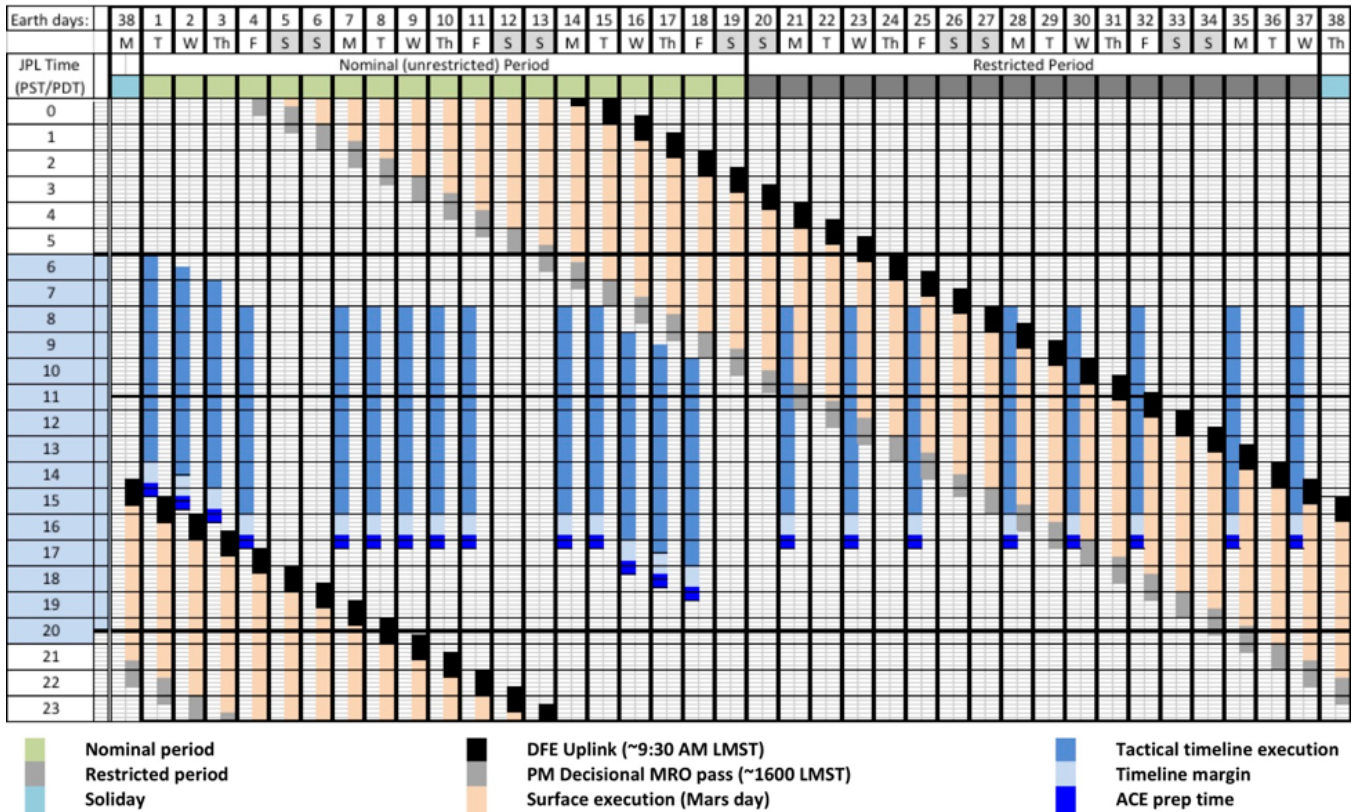


Figure 1. MSL tactical timeline with MRO+ODY.

2. BACKGROUND

MSL Tactical Planning Timeline Overview

Mars surface operations requires tactical planning in order to respond to the latest rover state before proceeding with the next set of science and engineering activities. Due to lighting and thermal constraints, MSL activities are primarily scheduled during the Martian day. This drives the timing of the uplink of the planned activities to the rover and the downlink of the resulting rover telemetry and data products to the operations team. The downlink data is termed “decisional”, because it is the rover data needed to inform the planning team’s decisions as to what activities to perform in the next plan. MSL relies on Mars orbiter relay between the rover and Earth in order to achieve the downlink timeliness and throughput required for operations. The time between receipt of the latest rover telemetry and the deadline to radiate the planned commands and sequences to the rover bounds the tactical timeline. Due to phasing between Earth and Mars, the timing of the downlink and uplink windows shift by about 40 minutes each day as depicted in Figure 1. Additionally, the availability of the Deep Space Network (DSN) antennae to radiate the uplink products and the latency for an orbiter to return the relayed rover data can impact the timeline.

In order to support a sustainable operations schedule that mitigates human fatigue, MSL performs tactical planning between 06:00 to 19:30 PT. As a result, there are “restricted” periods when the downlink is too late or the uplink is too early to enable tactical planning. Operations efficiency is the ratio between the number of unrestricted or “nominal” planning days to the number of Martian days (sols). The greater the operations efficiency, the more the operations team can effectively interact with the rover. For additional information,

please refer to Sharon Laubach, “Calculation of Operations Efficiency Factors for Mars Surface Missions” [8].

Relay Planning Overview

Mars relay planning consists of determining available orbiter overflights, identifying viable candidates for relay sessions, negotiating between the orbiter and lander teams, and sequencing the negotiated relay configuration. Relay planning generally takes place in two-week cycles called *planning periods*. The primary tool used to coordinate relay planning is the Mars Relay Operations Service (MaROS)[7]. Orbiter ephemeris data is used to populate MaROS with the times when the orbiter is in view of the lander as overflights. That information, along with the Orbiter Sequence of Events (OSOE), is used by the lander teams to request relay sessions. The OSOE captures information such as when the orbiter is unavailable to support relay and when DSN antenna tracks are allocated to support the downlink of relay data.

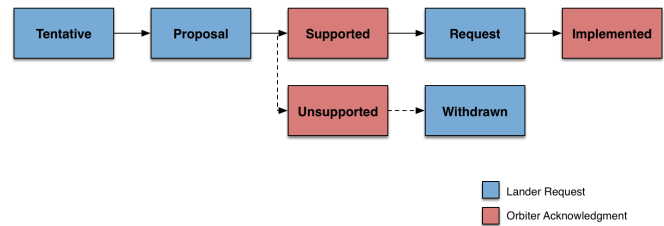


Figure 2. MaROS relay planning process.

Once the orbiter ephemeris and OSOE data is published to MaROS, the MSL relay planning team initially defines *tentative* requests to characterize the available level of relay

Rover Execution

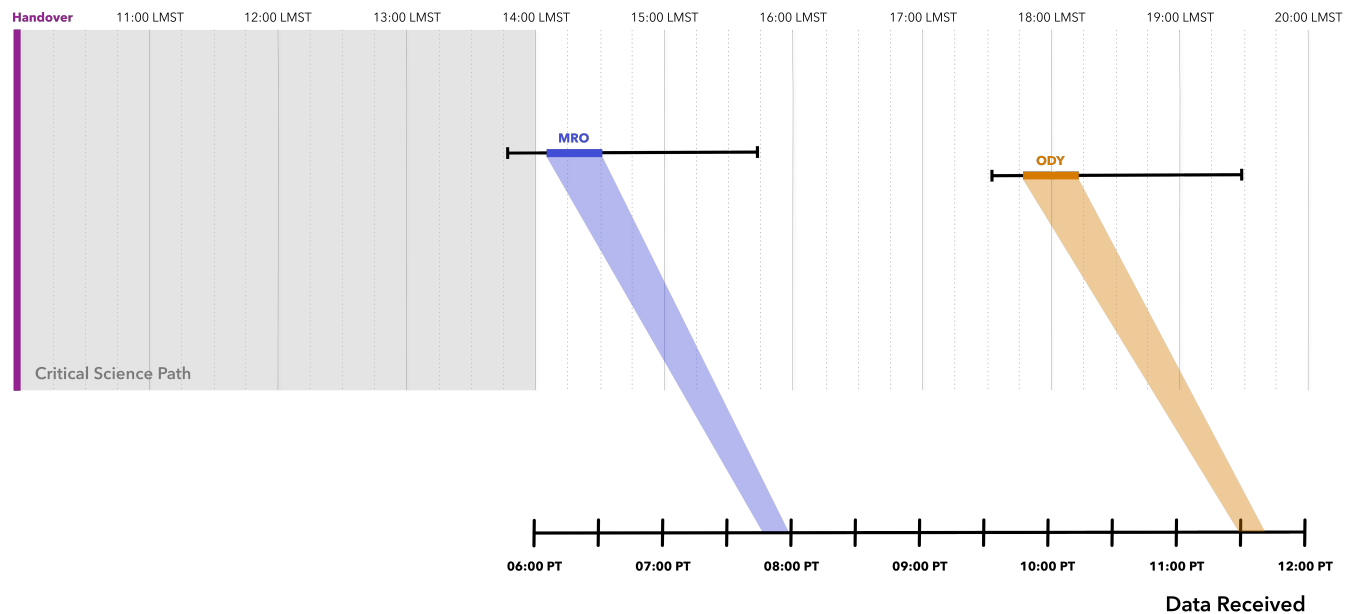


Figure 3. MRO and ODY overflight timing.

support to enable tactical planning. After the orbiter teams provide more refined ephemeris and OSOE data, MSL then publishes *proposal* requests. Proposals communicate to the orbiter team the relay sessions the lander team is requesting to schedule. The orbiter team then acknowledges the relay sessions as “supported” or “unsupported”. If unsupported, the lander team withdraws the request. Otherwise, the lander team finalizes the relay support configuration to be sequenced by both the lander and orbiter teams as the final *requests*. The due dates for proposals and requests for a given planning period cycle are captured in Figure 10 in the Appendix. Following requests, the orbiter and lander teams meet at the Short-Range Relay Coordination to approve the planning periods for implementation and execution.

Pre-InSight Era with MRO and ODY

Since MSL’s landing in 2012, MRO and ODY have been the primary relay assets used to downlink MSL’s data. Their sun-synchronous orbits result in temporally consistent overflights, which the initial MSL mission design relied upon to ensure sufficient decisional data return for tactical planning. As depicted in Figure 3, MRO and ODY overflights consistently occur in the Martian mid-afternoon and evening, respectively. This allows sufficient time after the uplink window at about 10:00 LMST (Local Mean Solar Time) to perform science observations before downlinking the collected data to enable tactical planning of the next sol. This time between the uplink window and the decisional pass is called the Critical Science Path (CSP). The temporal placement of MRO maximizes operations efficiency, since the timing is late enough to enable the CSP but early enough to return decisional data. ODY is often too late to be decisional, but can contribute to the total decisional data return when uplink is earlier in the planning day (see Figure 1).

Such consistency with MRO and ODY also lent itself to straightforward MSL relay planning. All MRO and ODY overflights that were considered “usable” were requested for relay.

Overflights are usable if they meet a minimum set of criteria:

- The maximum elevation must exceed 10 degrees. Anything less is unreliable and, therefore, is not a viable relay.
- The overflight cannot occur during any orbiter non-relay periods as captured in the OSOE.

The overflights that pass through those filters are then requested, resulting in about four overflights per sol: one MRO and one ODY PM overflight between 14:00-19:30 LMST, and one MRO and one ODY AM overflight between 02:00-07:30 LMST.

3. CHALLENGES WITH THE INTRODUCTION OF INSIGHT, MAVEN, & TGO

Impacts of InSight’s Arrival on Mars

With InSight’s landing site in close proximity to the MSL rover, the two missions now share orbiter view periods and thus overflight opportunities must be distributed (or potentially shared) between the two landers. In assessing how to implement this coordination, the MSL and InSight teams had to first determine the needs of both missions, and then determine how to incorporate new capabilities available from the Mars Relay Network orbiters to satisfy those needs.

As discussed in Section 2, MSL has historically relied on consistent MRO and ODY PM overflights to downlink decisional data and maximize operations efficiency given the staffing constraints of the MSL project. Similarly, InSight needs decisional data from the Martian PM passes during the first couple months while they are deploying instruments. Once InSight’s instruments are deployed, their decisional data needs are reduced, but they still need sufficient relay opportunities to downlink the science data collected by their instruments.

Additionally, MSL benefits from the high data volume return

associated with Adaptive Data Rate (ADR) capabilities provided by the MRO, MAVEN, and TGO orbiters. ADR allows the orbiter to monitor the signal-to-noise ratio (SNR) of the radio signal from MSL and update the data rates dynamically, up to 2048 kbps [4]. Conversely, InSight cannot use the ADR capability as they are equipped with a CE-505 radio [4]; therefore, all InSight overflights are sequenced in advance with fixed data rates up to 256 kbps. ADR provides MSL more data volume per relay opportunity for ADR-enabled orbiters, which can help offset MSL's greater decisional data volume needs.

Both MSL and InSight would also like to maintain operational capabilities with a diversity of orbital relay assets. This ensures that both missions are robust to any unforeseen circumstances that may reduce the downlink returned by any single orbiter.

Due to their proximity, MSL and InSight also have to manage any potential crosstalk interference. InSight and MSL will be the first landers that could be utilizing two different Mars Relay Network orbiters at the same time while both landers are in view of both orbiters. Because both landers will be using single-access proximity links, there is a risk that the orbiter hails could interfere with each other [4]. Since Proximity-1 protocol allows for re-hailing at any time during the overflight, there cannot be any overlap across two different overflights between the MSL and InSight projects. Crosstalk can have a significant impact on mission return if two decisional overflights conflict, since one mission would have to lose a decisional overflight.

As an effort to maintain operations efficiency for both landed missions despite the issues caused by crosstalk, the Mars Relay Network developed the capability to split an overflight. This "split-pass" or "split-relay" capability allows both landers to share a single relay opportunity that otherwise could have only been used by one lander [4]. This is done by splitting the total shared view period into two individual relay sessions (one for each lander).

The MSL and InSight teams collectively determined the best strategy for how to distribute the overflights from all Mars Relay Network assets while taking into consideration the needs, capabilities, and constraints of all missions. From these initial discussions, it became clear that sharing the data return capabilities of MRO and ODY between both landers would not be sufficient to meet MSL's operations efficiency needs, and that MSL would benefit from scheduling additional ADR-enabled MAVEN and TGO overflights to provide additional data return. However, relying on these assets introduces their own challenges.

Non-Sun-Synchronous Orbiters (MAVEN & TGO)

The addition of both MAVEN and TGO orbiters has necessitated a shift in paradigm from the consistency of the MRO and ODY paradigm that has existed since landing in 2012. While MAVEN and TGO both present their own individual challenges, both orbiters also occupy non-sun-synchronous orbits that precess, so their view periods to "walk" sol-to-sol. This characteristic directly affects the usability of MAVEN and TGO as decisional assets. As their orbits precess over time, daily MAVEN and TGO overflights may occur too early in the sol and conflict with the CSP or too late in the sol to be decisional for the subsequent planning cycle. This effect persists until view periods have shifted enough to once again allow for decisionally usable overflights. MAVEN and TGO's walking view periods also exacerbate the previously

mentioned challenge of crosstalk as they result in an increase in conflicting overflight opportunities for MSL and InSight.

While MAVEN and TGO present new sets of challenges, both orbiters also exercise great ADR performance and provide overall larger data return as a result. MAVEN's highly elliptical orbit enables overflights up to 30 minutes in duration; this is double the average overflight duration for MSL and thus increases data return. Additionally, TGO and MAVEN have taken significant efforts to reduce Electro-Magnetic Interference (EMI) on their UHF relay links as a result of the "relatively large 4 dB degradation in Electra threshold performance" observed in-flight for MRO due to two of MRO's science instruments [5]. Post-launch modifications to MRO's Electra firmware as well as shifting to operating in an EMI "quiet mode" during relay operations allowed MRO's EMI performance degradation to decrease from about 10 dB to about 3-4 dB [2]; whereas, MAVEN's observed EMI degradation on Electra performance is only about 0.7 dB as a result of increased attention to EMI during development and testing [3].

MAVEN's Orbit

The MAVEN orbiter poses many challenges by not only occupying a non-sun-synchronous orbit, but also a highly elliptical one with an apoapsis of 6,000km and a periapsis of 150 km that skims Mars' atmosphere [5]. The MAVEN spacecraft's orbit allows the mission's science team to achieve their science goals through the combination of providing coverage across all altitudes and precessing in latitude and local solar time, thus allowing for complete coverage of Mars' upper atmosphere [9].

As a result, MAVEN's view period durations have significant seasonal variability for MSL, ranging in durations from as short as 10 minutes to upwards of 2-3 hours. That being said, MAVEN relay sessions are limited to 30 minutes due to antenna thermal constraints for both spacecraft [3]. Since MAVEN view periods are typically longer than 30 minutes, this poses an interesting problem for the MSL planning team, who must select their preferred 30 minutes within a given view period. This decision is exacerbated by MAVEN's elliptical orbit. Elevation angle and range from the rover vary significantly in the view period, which can greatly affect overflight data return. Additionally, within a given 2-3 hour view period, data return latencies will vary as well. The various potential geometries and data return latencies must be weighed to determine the optimal 30-minute segment of a view period. MAVEN's orbit precession also means that these view period geometries will also shift, resulting in seasonal variation in data return. The MSL planning team's approach to determining which 30 minutes of these view periods is best will be discussed more in Section 4.

MAVEN's orbit also produces periods of challenging thermal and power constraints due to the frequency of eclipses. This effect reduces the total number of overflights that MAVEN can support, yielding fewer MSL relay opportunities.

A final challenge posed by MAVEN relay operations relates to orbit reliability and uncertainty. Due to its low periapsis and proximity to Mars' atmosphere, the spacecraft's orbit is affected by atmospheric drag. This effect is difficult to predict and thus creates a large unknown in the relay planning process for the MSL team. Consequently, data return predictions can be significantly impacted as planned overflight geometries (elevation angle and range) evolve. In the extreme case, the MAVEN spacecraft may even be below

the horizon for a portion of a previously planned overflight. Adding further complexity, the MAVEN team often executes Orbital Trim Maneuvers (OTMs) to adjust the spacecraft's orbit. OTMs are another cause of shifting geometries, which can significantly impact data return or even cause complete loss of a planned overflight. Frequently, OTMs are planned and executed too late in the relay planning process for the MSL team to react to.

TGO's Planning Process and Timeline

The challenges with using TGO for relay are largely due to the differences between TGO and the NASA orbiters' relay planning processes and timeline. While the relay planning process for the NASA orbiters begins approximately 2-3 weeks prior to when the overflights execute, the TGO planning process requires identifying relay opportunities and submitting proposals approximately 3 months in advance.

At 12 weeks prior to overflight execution, MSL must select and submit proposals for overflights that the team is interested in scheduling, reserving TGO's on-board resources for relay (e.g. timing, data volume). Because TGO and MAVEN overflights walk, they may conflict with each other or with MRO and ODY overflights. However, at this point in time, the orbital ephemeris predicts for the NASA orbiters are not well known. The DSN schedule has also not been finalized, so it is hard to accurately predict when data returned on specific overflights will arrive on Earth and be available for planning. This uncertainty in the orbital ephemerides and data latencies makes it difficult for MSL to fully assess overflight opportunities and make decisions as to which TGO overflights to propose 3 months in advance.

Any overflights that MSL selects and proposes at this stage of the process, and then later determines is no longer usable or desirable, must be withdrawn. This can negatively impact TGO's science and relay planning processes since the TGO science team cannot recover that unused time for science. To minimize disruptions to TGO science, it is important that MSL only proposes overflights intended to be used, despite the lack of information at the time. InSight must also propose TGO overflights at this time, which adds further complexity to the relay planning process since MSL and InSight must coordinate to prevent requesting conflicting TGO overflights. Any overflight that the two landers may want to share as a split-pass must be identified and marked as such at the time of proposal submission, as it is not possible per TGO's processes to later convert overflights proposed by a single lander to a split pass.

These challenges required the MSL and InSight missions to agree on how to allocate and potentially share TGO overflights so far in advance, while still enabling decisions regarding the NASA orbiters' overflights much closer to overflight execution.

InSight Negotiations

The InSight and MSL teams have arrived at an agreement for how to split overflight opportunities amongst both missions. This agreement defines a baseline relay utilization plan both for InSight's first few months of operations, when critical deployment activities and checkouts are scheduled as well as for ongoing nominal operations when InSight's primary science observations will be collected.

In general, the two missions divided the sun-synchronous MRO and ODY overflight opportunities evenly, where each

mission is allocated one PM and one AM overflight per sol. Which mission is allocated which orbiter depends on both InSight's mission phase (i.e. the deployment phase or nominal science operations) and InSight's tactical timeline (i.e. how the phasing of the Earth and Mars days affects the planning schedule on Earth), which impacts how much decisional data is needed and when it must arrive on Earth. MRO and ODY, with their sun-synchronous orbits, are relied upon by both MSL and InSight for their consistent timing with quick and reliable return of data to Earth. This baseline agreement helps satisfy the requisite operations efficiencies for InSight, but MSL's decisional data needs are still in excess of the support provided by MRO and ODY.

TGO and MAVEN overflights are also distributed amongst both landers. MSL and InSight are collectively allowed two TGO overflights per sol and up to one MAVEN overflight per sol, where a split-pass is only considered to be one overflight. TGO and MAVEN can only sometimes return decisional data due to their non-sun-synchronous orbits; however, they are also useful to both landers for non-decisional data return, especially for MSL due to the use of ADR. Because of MSL's higher decisional data volume needs and the benefits of ADR, both missions agreed that MSL has first selection for up to two TGO or MAVEN overflights per sol, with the intention to share one of those overflights as a split-pass with InSight. InSight then has the opportunity to schedule one of the remaining MAVEN or TGO overflights. With this agreement, both missions should satisfy their operations efficiency requirements while also satisfying their total data downlink requirements.

The only remaining issue to be resolved was the priority of overflight de-confliction when a crosstalk problem was identified in planning. Both missions agreed that in the event of crosstalk, the mission that was allocated the sun-synchronous orbiter would keep their overflight. This ensured that the decisional downlink needs are protected for both missions since: 1) the 2 sun-synchronous orbiters can never conflict because of their orbital geometry; and 2) the sun-synchronous orbiters regularly have a PM pass that could enable decisional data. In the event that a TGO and MAVEN overflight conflict between the two lander missions, both teams agreed that the TGO overflight would take precedence. This is because TGO overflights are planned much further in advance, so the conflict is reconciled during the MAVEN overflight selection process to ensure minimal impact on the science return of either orbiter.

With these agreements in place with InSight, MAVEN, and TGO, the MSL team had the requisite knowledge to design and develop a robust system of processes and tools to routinely and effectively coordinate, prioritize, and select the UHF relay sessions that would maintain MSL mission operations efficiency.

4. PROCESS AND TOOL IMPLEMENTATION

In order to enable the new paradigm and mitigate the aforementioned challenges, the MSL team redesigned and redeveloped their relay planning process and tool implementation. This primarily consisted of introducing MAVEN and TGO as nominal relay assets, enabling negotiation and coordination with InSight, and developing new overflight selection criteria to meet operational efficiency needs. The following sections describe the implementation approach MSL took to accommodate these changes.

Software Architecture

For non-trivial software development projects, software engineers often spend time at the beginning of a project defining high level system requirements and designing the most critical code structures or software architecture. Automating the process of selecting and requesting optimal overflights for relay is a complex problem with many interdependencies and constraints. Additionally, the tools have to support the evolving Mars Relay Network and MSL mission priorities. Therefore, the MSL development team spent the first weeks of the project defining the critical requirements and designing a supporting software architecture for the *create_uhf_requests* tool.

The purpose of the *create_uhf_requests* tool is to automate the overflight selection process. While a full software requirements specification is outside the scope of this paper, this section informally summarizes the requirements that drove the team’s architectural decisions:

Functional Requirements

- Automate the overflight request process using a configurable strategy for selecting overflights. The complete strategy had not been defined prior to the start of this development project.
- Support interactive and fully automated modes of execution. For example: a relay planner executes and interacts with the script as it selects and requests overflights or a schedule task that executes without human interaction.

Non-Functional Requirements

- A domain expert must be able to create, update, and remove overflight selection rules as the priorities of the project change.
- Solution must be able to adapt to changing operating environments with minimal and isolated code changes.
- MSL Developers must be able to test and evaluate the results of applying overflight selection rules against various scenarios.
- MSL Developers must be able to automatically test interfaces to external systems. As the operating environment evolves, developers must be able to quickly test proposed changes to external services.

Constraints

- Must operate in MSL’s current operations environment and use technology that is familiar to MSL developers.
- Short schedule with hard deadline of InSight landing

Interfaces

An important initial step in designing a tool is to explore how it will fit into the existing environment and define the interfaces the tool must provide and depend on. A context diagram is a useful tool for driving software requirement discussions. The MSL team created Figure 4 while defining requirements.

Based on these requirements, the MSL development team designed a software architecture using several design tactics. In the field of software architecture, a design tactic represents a specific code structure, technique, or pattern to achieve a desired quality in the software [1]. Table 2 summarizes the tactics used by the development team and is followed by a detailed explanation of how each tactic was applied.

First, the MSL development team had to decide how to break this project into logical parts and assign development responsibility. Team composition, experience, and skillsets

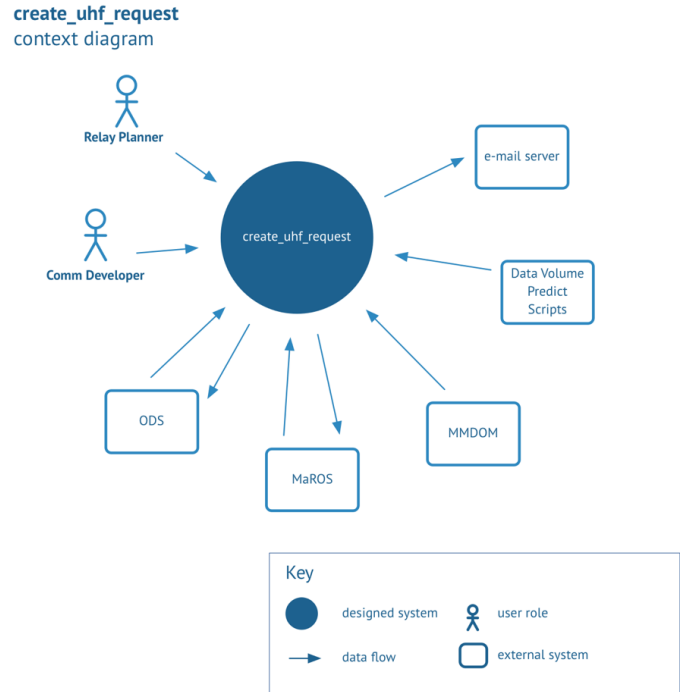


Figure 4. *create_uhf_request* context diagram.

Table 1. Interfaces.

Interface	Description
Relay Planner	Comm team member responsible for overflight selection request
Comm Developer	Software engineer with Comm domain expertise
MaROS	Mars Relay Operations Service, a centralized service used by MSL, InSight, and Mars orbiters to coordinate relay planning
ODS	MSL file storage for operational and historical data
MMDOM	Provides access to mars orbiter specific information not available through MaROS
Data Volume Predict Scripts	Generates data volume predictions for orbiter overflights

Table 2. Design tactics.

Tactic	Requirement
Decompose solution into domain specific and infrastructure specific modules	Team consists of three part-time individuals: 1. Domain Software Engineer 2. Infrastructure Software Engineer 3. Software Architect
“Plug-in” interface for configuring overflight selection and modification	Automate the overflight request process using a defined but configurable strategy for selecting overflights
Adapters for each external system dependency	Adapt to changing environment with minimal and isolated code changes

where factors that the team considered and lead to the decision to divide the project into domain and infrastructure specific modules. Table 3 describes the code structure and implementation responsibility.

Table 3. Code structure and implementation responsibility.

Module	Description	Implementer
Config	Tool settings	Infrastructure Software Engineer
Domain	Domain specific logic Controllers: implement workflow Models: data structures and validation	Domain Software Engineer
Services	Communicate with external systems	Infrastructure Software Engineer
UI	Command line UI interaction	Infrastructure Software Engineer

After the team defined the top-level code structure, they focused on how to meet the overflight selection requirements. Detailed overflight selection rules were not completely defined at the start of the software project and the MSL development team knew that the requirements would likely change as the MSL mission evolves. The development team chose a “plug-in” structure that accepts an ordered list of selection rules; each rule module implements the “rule” interface. A comm developer can easily write, test, and add new rules as the mission’s priorities change. Using a configuration file, the comm developer specifies which rules to apply and in what order.

create_uhf_requests

Component Diagram - rules interface

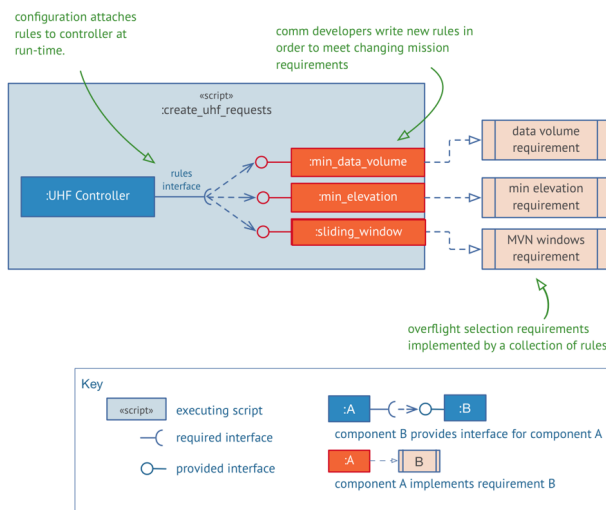


Figure 5. create_uhf_requests component diagram rules interfaces.

The rules plug-in tactic allows the tool to adapt to changing overflight selection criteria but does nothing to shield it from

changes to the operations infrastructure. To address this issue, the development team chose the adapter pattern. An adapter is a module of code that converts the interface of an external system to an interface that our tool expects [6]. Using the context diagram as a guide, the development team defined an interface for each external system requirement. For each interface, the team implemented an adapter that interacts with the external system. With this structure, developers can simply and quickly test each adapter in isolation. When proposed changes to an external service are announced, the team can run a suite of automated tests on the specific adapter to proactively identify issues.

Figure 6 shows an example of the create_uhf_requests tool configured to use a new data volume prediction tool. In order to use the new data volume prediction system, a developer writes an adapter, then updates the tool configuration to point to the new implementation (advanced_data_vol). Developers do not need to modify the code or logic inside the UHF Controller or other domain packages to take advantage of the new service.

create_uhf_requests

Component Diagram - service interfaces

Developers swap service implementations through a configuration file without touching the domain logic in the controller

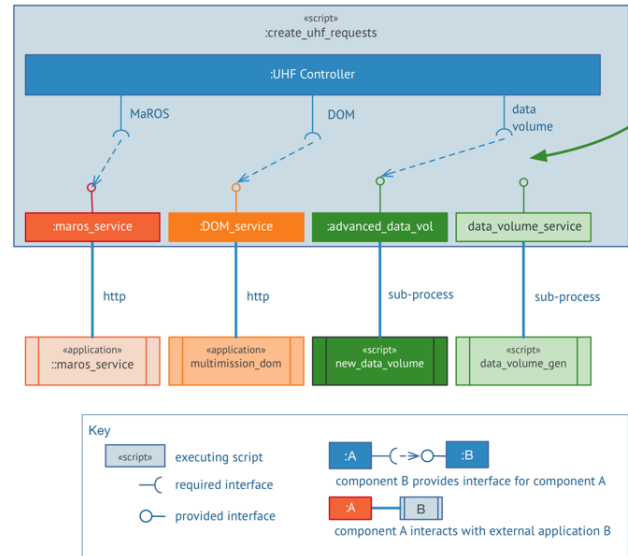


Figure 6. create_uhf_requests component diagram service interfaces.

At the start of the development project, The MSL team invested precious time to define the most critical requirements and to design a supporting software architecture. The result is an organized, configurable, and adaptable tool that will continue to meet the needs of the MSL’s overflight selection process for the remainder of the mission. Furthermore, the architecture and the components built to support overflight selection can be used to guide and construct the remaining set of tools to support MSL relay planning and provide a blueprint for future Mars missions.

Rover Execution

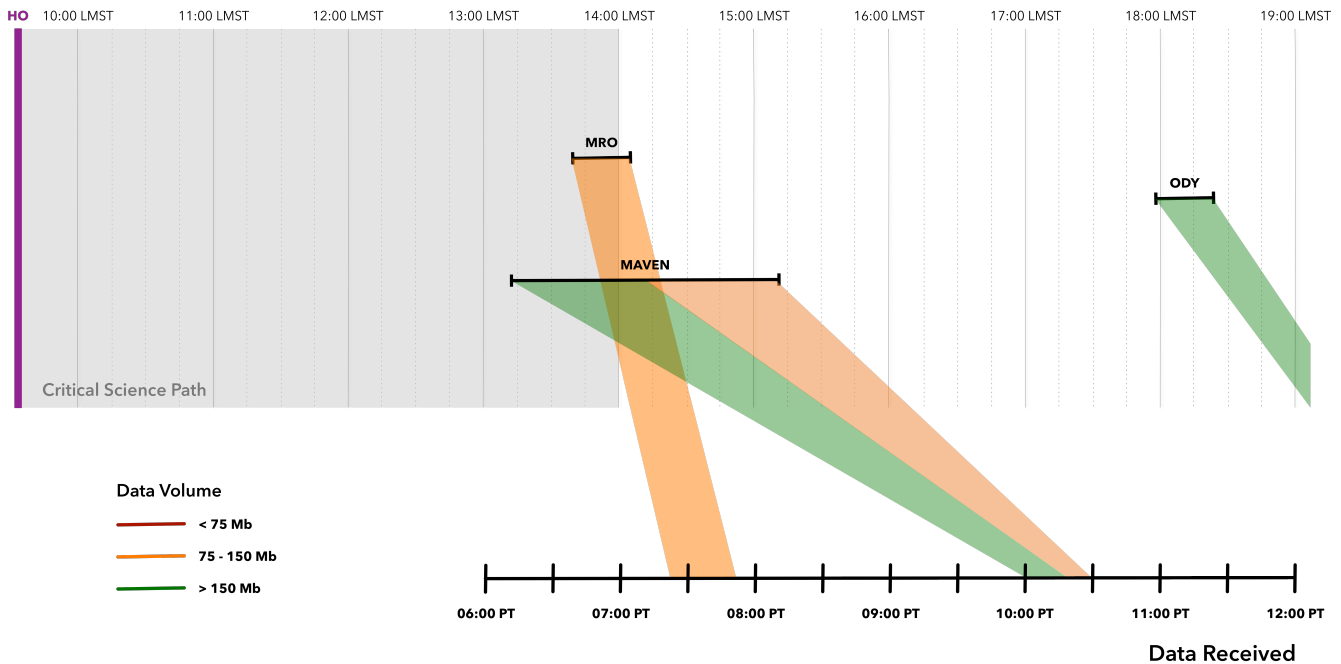


Figure 7. Decisional use case.

Overflight Selection

As discussed in Section 2, overflight selection for MRO and ODY consisted of requesting all “usable” overflights that meet minimum criteria. Overflight usability for MAVEN and TGO requires additional considerations since the overflights can walk and conflict with the uplink window and Critical Science Path (CSP). Additionally, as discussed in Section 3, one PM MRO or ODY overflight and one AM MRO or ODY overflight are allocated to InSight depending on InSight’s mission phase and tactical timeline. Thus, overflights must meet the following criteria to be considered usable in the new paradigm:

- The maximum elevation must exceed 10 degrees and the predicted data volume must exceed 20 Mb. Anything less is unreliable and, therefore, is not a viable relay.
- The overflight cannot occur during any orbiter non-relay periods as captured in the OSOE.
- The overflight cannot conflict with the uplink window or with the minimum CSP (2.5 hours after uplink).
- The overflight cannot be allocated to InSight based on InSight’s mission phase and tactical timeline.
- MAVEN and TGO overflights that occur during InSight’s MRO and ODY allocation must deconflict with InSight’s overflight timing due to crosstalk.

The overflights that pass through those filters then need to be further down-selected. Due to MSL’s resource constraints, up to 4 overflights are selected per sol. To maximize MSL’s operations efficiency, the overflight down-selection prioritizes decisional overflight selection followed by additional data volume return.

An overflight is considered decisional if the data is down-linked in time to plan: either by the latest planning shift start time (11:30 PT) or by the uplink deadline minus the tactical timeline duration (8 hours), whichever is earlier. With multiple decisional overflights potentially available, it

is not always obvious which decisional overflight is best. For example, in Figure 7, both an MRO and MAVEN overflight are decisional. The MRO overflight is towards the end of the CSP, and its data is down in time for an 8:00 PT shift start, but it only returns 80 Mb. The MAVEN overflight is 2-hour long, but only up to 30 minutes can be used for relay, as discussed in Section 3. The first half has high data return (≥ 200 Mb), but its timing conflicts with the CSP. The second half has mediocre data return (75-100 Mb), but its timing enables more science. Additionally, the data arrival time of the MAVEN overflight results in a 10:30 PT shift start time, which negatively impacts MSL’s planning and science team that participates from all over the world. The uncertainty of the MAVEN orbit as discussed in Section 3 is also a consideration as to whether or not to rely on it for decisional data.

In order to select from available decisional overflights, MSL came up with decisional selection criteria to prioritize overflights based on three metrics:

1. The duration after the uplink to enable the CSP
2. The total data volume (DV) returned
3. The resulting tactical shift start time based on the data arrival time

Quantitative thresholds for each metric defined based on mission and science priorities are used to populate the decisional selection filter table shown in Table 4. All “usable” overflights then pass through the filter table in priority order. If one or more overflights meet a given priority’s criteria, the tiebreaker is used to select the decisional overflight. This table affords the mission flexibility in prioritizing the metrics relative to each other, while also enabling implementation and configuration as mission priorities evolve. In cases where the first selected decisional overflight does not have sufficient data volume to enable planning, another decisional overflight is selected using the filter table.

Table 4. Decisional selection filter table.

Priority	CSP (hrs)	\geq DV (Mb)	Shift Start (\leq hrs from 08:00 PT)	Tiebreaker
1	6:00	250	1.5	Orbiter
2	6:00	120	1.5	Data Volume
3	5:15	250	1.5	Orbiter
4	5:15	120	1.5	Data Volume
5	6:00	80	1.5	Data Volume
6	5:15	80	1.5	Data Volume
7	4:30	80	1.5	Data Volume
8	6:00	50	1.5	Data Volume
9	5:15	50	1.5	Data Volume
10	4:30	50	1.5	Data Volume
11	4:30	80	3.5	Shift Start
12	4:30	50	3.5	Shift Start

After one or two decisional overflights are selected, the remaining overflights are selected based on data volume return. The overflights with the greatest data volume that occur after the decisional pass(es) are selected until the goal total data volume per sol is met or max four overflights per sol are selected. Due to the higher data volume return achieved with MAVEN and TGO, MSL's total data volume goal can be met with fewer than the historical four overflights per sol, releasing duration and energy for more MSL science.

MAVEN Segments

In order to select the best 30-minute window of a MAVEN view period, as described in Section 3, the view period is discretized into 30-minute segments offset by 1-minute increments as depicted in Figure 8. Each segment is treated individually when passing through the decisional selection filter table since each segment could have different metrics. Once a segment of a given MAVEN overflight is selected, the remaining segments are removed from consideration.

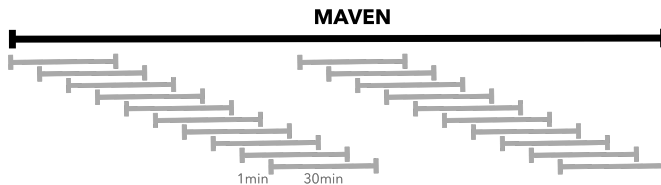


Figure 8. MAVEN segments.

TGO Proposals

As discussed in Section 3, TGO proposals are due 12 weeks prior to execution. Due to the uncertainty of other orbiters' overflights at this time, MSL is unable to make informed decisions as to which overflights to select. Additionally, only 2 TGO overflights per sol are allocated to MSL and InSight.

Therefore, both missions must coordinate as to which TGO overflights are proposed.

In order to provide the most flexibility, while also minimizing the impact to TGO, MSL and InSight split two TGO overflights per sol at TGO proposals. This way, at requests, the TGO overflights can remain split or be completely allocated to one mission or the other without exceeding the timing of the proposed relay session. MSL uses the selection logic described in Section 4 to select up to two TGO overflights per sol. The other orbiters are not considered during selection due to the uncertainty captured in Section 3. InSight then splits any MSL-selected overflights that InSight deems usable. Additionally, if MSL does not select a second overflight, InSight has the opportunity to select the second based on their criteria.

InSight Coordination for MRO and MAVEN Proposals

Prior to MRO and MAVEN proposals (refer to Appendix), MSL and InSight must negotiate and coordinate which overflights are allocated to which mission. MSL and InSight have the opportunity to split shared view periods or allocate the overflight to a single lander. Any other overflights in view need to be deconflicted to prevent crosstalk.

As discussed in Section 3, each mission is allocated one PM MRO/ODY overflight and one AM MRO/ODY overflight. Because the MRO and ODY allocation has precedence, MSL begins by blocking out InSight's allocated MRO/ODY overflight to ensure not only that MSL's corresponding MRO/ODY overflight isn't selected but also that all other MAVEN/TGO overflights deconflict for crosstalk. If any overlapping overflight only partially conflicts with InSight's overflight, the usable portion is considered. Then, MSL uses the logic described in Section 4 to select overflights.

Between MSL and InSight, only 1 MAVEN and 2 TGO overflights can be requested per sol. MSL can only select one of each or 2 TGO overflights per sol. The remaining TGO or MAVEN overflight is allocated to InSight. Therefore, once the algorithm selects 1 MAVEN or the second TGO overflight, the remaining MAVEN/TGO overflights are excluded from consideration.

Due to the offset in orbiter planning period boundaries, OSOE information is missing when the planning periods do not overlap (see Appendix). However, MSL needs to be able to make informed decisions as to which overflights to select in order to enable InSight coordination. To mitigate this, MSL currently assumes a 3-hour latency for all MRO and ODY overflights and ignores all MAVEN overflights without OSOE information. Once MAVEN latencies are better characterized, considering MAVEN overflights without OSOE information can be reassessed. Additionally, MSL has the opportunity to withdraw unnecessary overflights at requests when the missing OSOE information is available.

After MSL submits their selected overflights as tentative, InSight reviews their opportunities to ensure they are sufficient to meet their needs. MSL and InSight then bring any deviations from the agreed upon allocation to the MSL-InSight Coordination Meeting to be negotiated prior to submission of proposals.

Requests

At the time of requests, more reliable ephemeris and OSOE information is introduced. Additionally, preliminary OSOE information is available for the subsequent planning period

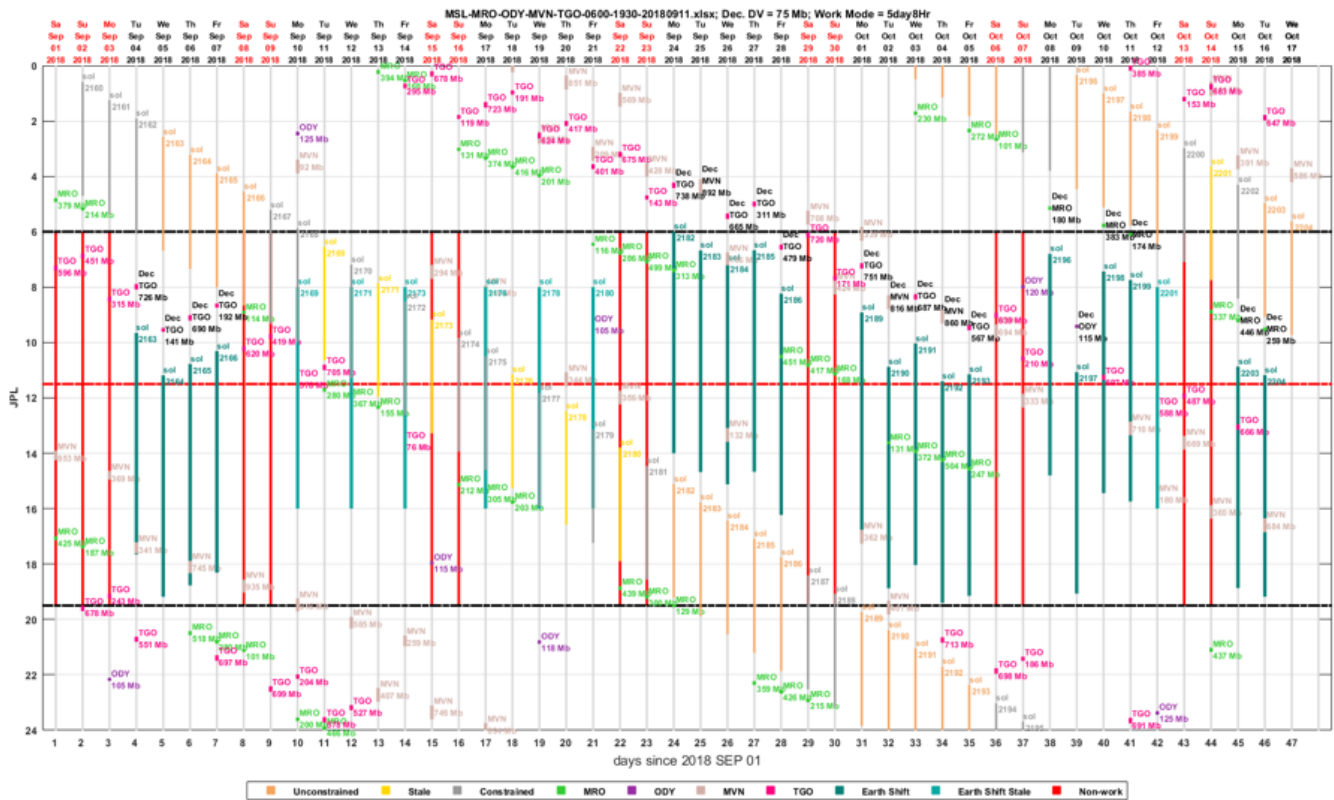


Figure 9. MSL tactical timeline with MRO+ODY+MAVEN+TGO.

to handle the planning period boundary condition. Requests give the lander teams the opportunity to respond to the latest information and update the overflight timing and parameters as long as the requests are within the proposed allocation. Any deviation from the InSight coordination is negotiated ad-hoc.

5. PLANNING TIMELINE IMPACTS

MSL operations to date has been used to a 38-day tactical timeline cadence, reflective of the relative phasing of Earth-time and Mars-time (which also uses a 24-hour clock, where a Martian hour is approximately 1.02749 Earth hours long), as illustrated in Figure 1. This cadence comprises a period of roughly 2.5 weeks of “nominal” planning days, followed by a period of roughly 2.5 weeks of “restricted” planning days, and finally a “soliday”, a day without a planning cycle to bring Earth- and Mars-times back into synchronization.

“Nominal” days are those for which the Earth-based operations team has in-hand telemetry from the immediately preceding sol and can complete a planning cycle within its operational window (the period of an Earth day in which a tactical operations shift can be scheduled) and before the commands must be radiated to the rover the (Mars) morning of the sol being planned. During “restricted” days, either the telemetry cannot arrive in time before the planning shift must start, so the team must use stale data to plan rover activities, or the planning shift cannot complete in time to meet the deadline for radiation. The term “restricted” refers to the fact that under these conditions, the use of rover mobility and arm commands are carefully restricted to avoid planning

further terrain interaction without having current knowledge of the results of prior commanded interaction. This reliably repeatable cadence is enabled by the exclusive use of sun-synchronous relay assets, MRO and ODY, to return MSL telemetry.

In turn, this cadence results in an easily predictable staffing schedule for the operations teams. As shown in Figure 1, the cycle begins with a couple of days at the earliest allowable start time (6:00 AM Pacific), followed by a stretch of roughly 1.5 weeks at the “standard” start time (fixed at 8:00 AM Pacific), finishing the “nominal” period with a few “late slide” days, whose start times are pushed later by the telemetry arrival time until the latest allowable start time is reached (11:30 AM Pacific). During the “restricted” period, the standard start time is used—and further, due to the issue of stale data, planning is only done every other day (Monday, Wednesday, and Friday). Both engineering personnel located at the Jet Propulsion Laboratory and science team members located around the world can plan weeks and months in advance for staffing needs, knowing that only small perturbations to this cadence will occur.

On the face of it, the introduction of non-sun-synchronous relay assets, MAVEN and TGO, has the potential to disrupt this cadence due to the fact that their overflight times “walk” through the sol, as described in Section 3. In addition, losing the use of one of the two sun-synchronous orbiters for MSL data, due to the necessity of sharing relay assets with InSight, reduces the timeline cadence’s robustness to DSN downtimes, orbiter safing, and other changes to expected data return latencies from the brace of sun-synchronous relay assets. Work is ongoing to assess the full extent of the

- The addition of MAVEN and TGO to a robust sun-synchronous relay cadence tends to improve the overall ops efficiency, by (when the phasing of the non-sun-synchronous asset is advantageous) lengthening the nominal period of the cadence and/or interrupting the restricted period with additional nominal planning days.
- However, care must be taken to select decisional passes which do not introduce “whiplash” in the tactical shift start time—that is, which do not allow the start time of a given day to exceed 1.5 hours later than the prior day’s start time, nor 0.5 hours earlier than the previous day’s start time. Human factors fatigue countermeasures contraindicate excessive shift start time changes for individual team members, who are often scheduled for several planning shifts in a row.
- Further, the use of non-sun-synchronous assets for decisional downlink will result in a non-repeatable cadence, due to the different “walk” rates for MAVEN and TGO, and how those different rates beat against the 38-day Earth-/Mars-time phasing. The effects of these changes are yet to be assessed for the magnitude of their impacts on the ability to assess longer-time-horizon staffing needs for engineering and science team members.

As these impacts are further assessed and defined, MSL may then add new rules to the pass selection logic to: 1) reduce negative human factors effects such as whiplash, and 2) better leverage the additional relay support from MAVEN and TGO that seasonally allows MSL operations efficiency to increase.

The MSL team is expected to maintain its historical operations efficiency in the era of NSY arrival at Mars. To enable this, the team had to overcome several key challenges including: coordinating the usage of relay overflights with NSY to satisfy both mission's needs and properly deconflict crosstalk interference; understanding the uncertainties associated with the MAVEN spacecraft orbit; coordinating vastly different planning timelines between the tactically planned MSL and strategically planned TGO missions; and adapting the MSL mission operations staffing schedule to better accommodate the non-sun-synchronous orbiters. Through the use of modern systems engineering principles and software architecture design, the team was able to capture these various competing constraints and then design an automated suite of tools that will prioritize and select from the available overflights on each planning sol. This suite of tools satisfies the constraints of the interfacing missions, while also remaining flexible to MSL mission needs through the use of a configurable design. With the usable and selection filter capabilities, the MSL mission can deftly change their relay pass selection criteria without an overhaul of any processes or tools.

the overall MSL tactical timeline, resulting in a shift from an easily predictable operations staffing schedule (stemming from exclusive use of sun-synchronous orbiters for decisional data) to one that tends to improve overall operations efficiency and enable sufficient data return for both MSL and InSight but also may cause inconsistent planning start times and more unpredictable staffing schedules. Work is ongoing to assess the full extent of the impact of these changes, and overflight selection criteria may be updated to reduce negative human factors effects and make better use of the non-sun-synchronous orbiters to increase operational efficiency. The team will continue to monitor and adapt the solution proposed and may, eventually, be able to improve the total mission operations efficiency while also sharing overflights with NSY.

The problems and associated solutions discussed in this paper are relevant for future exploration missions on the surface of Mars and other planets where surface assets relay data through a network of orbiters. Additionally, the tools could easily be ported to current missions on the surface of Mars by adapting to each mission's constraints, planning processes, and mission needs. However, the constraints, challenges, and solutions captured in this paper could be used to inform the design and foundation of future space-exploration relay networks.

[illegible]

ACKNOWLEDGMENTS

The authors would also like to acknowledge Charles Lee,

whose integration of ops efficiency modeling with the Telecom Orbit Analysis and Simulation Tool (TOAST) aided with tactical timeline simulation.

REFERENCES

- [1] Bass, Len, et al. *Software Architecture in Practice*. Addison-Wesley Professional, 2012.
- [2] Bell, David, et al. "MRO Relay Telecom Support of Mars Science Laboratory Surface Operations." *Jet Propulsion Laboratory, California Institute of Technology*, IEEE, 2014.
- [3] Chamberlain, Neil, et al. "MAVEN Relay Operations." *Jet Propulsion Laboratory, California Institute of Technology*, IEEE, 2014.
- [4] Edwards, Charles D., et al. "Assessment of potential Mars relay network enhancements." *Jet Propulsion Laboratory, California Institute of Technology*, IEEE, 2018.
- [5] Edwards, Charles D., et al. "Relay Support for the Mars Science Laboratory and the Coming Decade of Mars Relay Network Evolution." *Jet Propulsion Laboratory, California Institute of Technology*, IEEE, 2012.
- [6] Gamma, Erich, et al. *Design Patterns, Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, 1994.
- [7] Hy, Franklin, et al. "Implementing Strategic Planning Capabilities within the Mars Relay Operations Service." *Jet Propulsion Laboratory, California Institute of Technology*, AIAA, 2011.
- [8] Laubach, Sharon. "Calculation of Operations Efficiency Factors for Mars Surface Missions." *Jet Propulsion Laboratory, California Institute of Technology*, AIAA, 2014.
- [9] "Science Orbit." MAVEN, Laboratory for Atmospheric and Space Physics, 2018, lasp.colorado.edu/home/maven/science/science-orbit/.